BNL 51976
UC-80
(General Reactor Technology — TIC-4500)

BNL--51976

DE87 009374

Probabilistic Risk Assessment

Robert A. Bari

Number 219

November 13, 1985

BROOKHAVEN NATIONAL LABORATORY
Associated Universities, Inc.
Under Contract No. DE-AC02-76CH00016 with the
United States Department of Energy

Printed in the United States of America
Available from:
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161

NTIS price codes:
Printed Copy: A02; Microfiche Copy: A01

PROBABILISTIC RISK ASSESSMENT

Robert A. Bari

The title of this talk, Probabilistic Risk Assessment, describes a methodology for analyzing the safety of nuclear power plants.

At the outset, I will tell you what this lecture is not about, because there are many important issues associated with nuclear power, and it is very easy to stray onto a larger plane. I will not discuss: economics of nuclear energy; safeguarding of nuclear materials; mining and fabrication; and finally, waste management. I will focus on nuclear safety and risk assessment. Starting out with a primer on nuclear power plants, I will talk about probabilistic risk assessment and provide a historical overview, and then tell you where we are now, concluding with the work that we are doing at Brookhaven Laboratory in the Engineering and Risk Assessment Programs. Finally, I will speculate about where we may be going.

Let us look at the sample space for these studies. Currently, 93 plants are licensed to operate in the United States, and, next year, the Nuclear Regulatory Commission expects that an additional 15 plants will be licensed, or will be close to being licensed. Then, for the long term into the early 1990s, another 15 plants are in the licensing pipeline. Basically this will be all. No further orders for plants have been placed by utilities, so the sample space is 123 nuclear power plants. These plants are scattered around the United States, with heavy concentrations in the Northeast, in the South, and also in the Midwest. There are some good reasons for this distribution. For example, in the Mountain States, other resources are available for producing electric power. The situation with the commercial plants in the United States is the following: except for the Fort St. Vrain Plant, the plants that generate electricity are all light water reactors. Light water refers to hydrogen, as opposed to deuterium, in the water molecules, and of these light water reactors, there are two basic kinds. One is the boiling water reactor (BWR), and the other is the pressurized water reactor (PWR). General Electric is the vendor for the BWR, as it is known. The PWR basically comes out of the Submarine Program, and Westinghouse is the main vendor, although Combustion Engineering and Babcock and Wilcox also supply PWRs. A large modern-size nuclear reactor generates about 1000 MW of electricity. Some of the older, and smaller reactors generate on the order of 400 to 600 MW. The basic flow chart from the fission process to the production of electricity is indicated in a simple "See Spot Run" fashion in Figure 1. We start with the nuclear fission process in the reactor, in which we basically get the kinetic energy from the recoiling fission products of ^{235}U. We get about

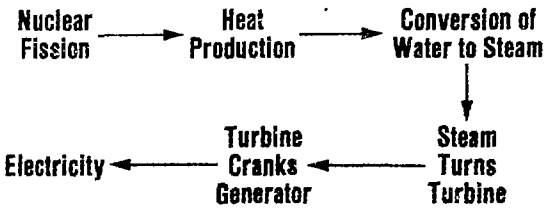

Figure 1. Nuclear fission energy conversion to electrical energy.

170 MeV from the fission products themselves. This leads to heating, as a result of the random thermal motion in the fuel matrix. Then, the heat is basically converted into steam, as a result of water being introduced into the reactor. This steam turns the turbine, which cranks an electric generator, to produce electricity. In a very general sense, this part of the cycle is similar to methods of generating electricity from fossil fuel plants. The technology is vastly different in the heat-generating part of the nuclear plants, but from that point onward the fundamental processes are the same for the steam energy transport and its conversion to kinetic energy in the turbine. The boiling water reactor is shown in Figure 2. The fission process occurs in the reactor core (not shown to scale), which is 12 feet high in the BWR. Water is brought into the core, boiled off into steam, which turns the turbine, which cranks the generator, to generate electricity. The schematized array shows the elaborate system of pumps and valves that come into the reactor. A

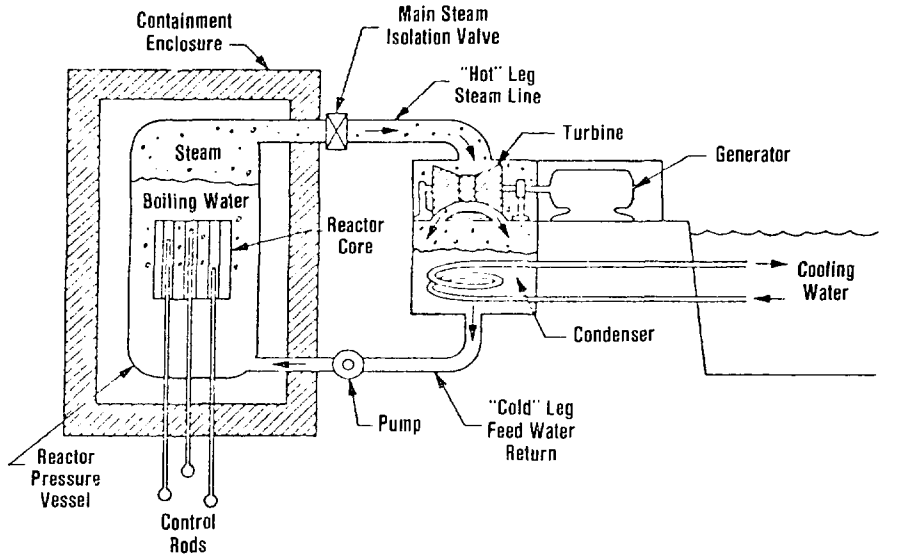

BWR

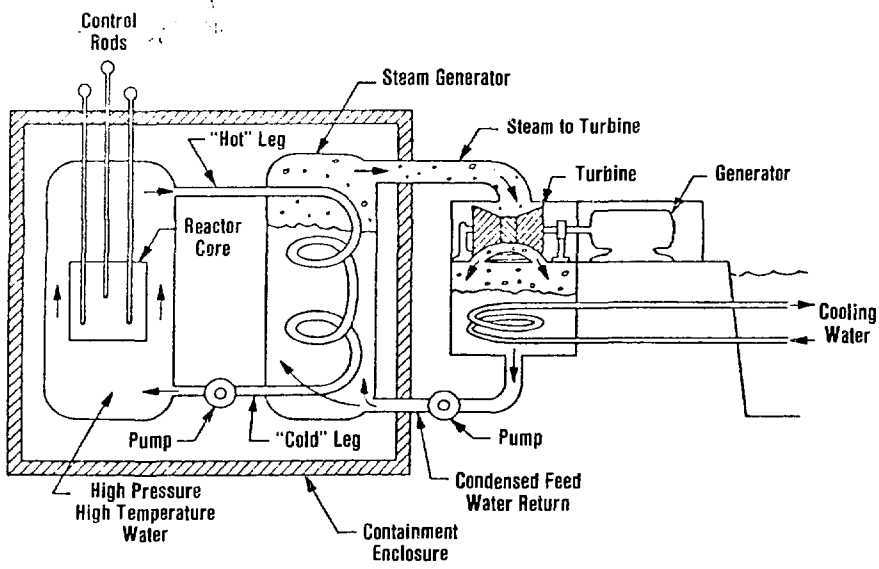

PWR

Figure 2. Schematic of a boiling water reactor (BWR) and a pressurized water reactor (PWR). Note the secondary loop in the PWR.

containment building surrounds the reactor to ensure that, in the event of some fission product release, the products are contained within the building. Around this there is another building, the reactor building; the turbine building encloses that part of the reactor.

The pressurized water reactor also shown in Figure 2 has an intermediate step. It does not boil water in the reactor vessel but in a steam generator: there is a secondary loop resembling the single loop of the BWR. The water, which is heated to about 600°F, is at about 2000 pounds per square inch in the vessel, while in the BWR case, it is at about 1000 pounds per square inch. The temperatures on the outlet of the core for both reactors are on the order of 500-600°F.

Figure 3 is a more elaborate view of a boiling water reactor vessel. Water is fed in for recirculation through an arrangement of jet-pumps, and then the water goes up through the core and steam is produced; there is a dryer which produces high-quality steam, then, finally, steam comes out of the vessel. Figure 4 gives a more detailed view of the containment building for the boiling water reactor. The vessel, which is about 50 feet high, is drawn to give a perspective on height, and the figure also shows the containment building and the so-called suppression pool. Should a pipe break in the reactor, which would lead to steam production and possible pressurization of the containment building, the arrangement is such that the steam is channeled down into the suppression pool, that is filled with ordinary water and acts as a pressure-reduction mechanism. Basically, steam is condensed in a very large suppression pool, which has a torroidal configuration and encircles the containment. Figure 5 illustrates a typical pressurized water reactor containment building with the reactor vessel. There is a much larger containment volume of about 2-1/2 million cubic feet. The steam generators are larger than the reactor vessel. The containment building, again, is a structure of reinforced concrete with a large base mat of about 9 feet.

Why are we worried about nuclear reactors? The hazard from nuclear power is in the fission products that are wrapped up in the core. The core of a typical large power reactor contains about a billion curies of radioactivity, which we would not want to be released, since high

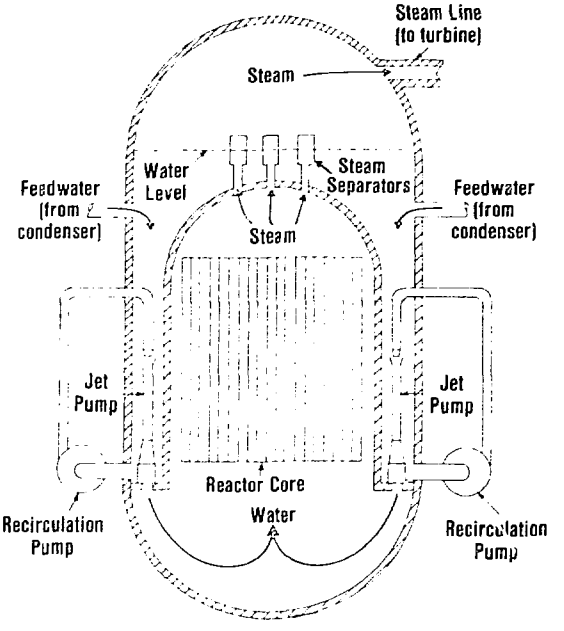

Figure 3. Details of the boiling water reactor vessel.

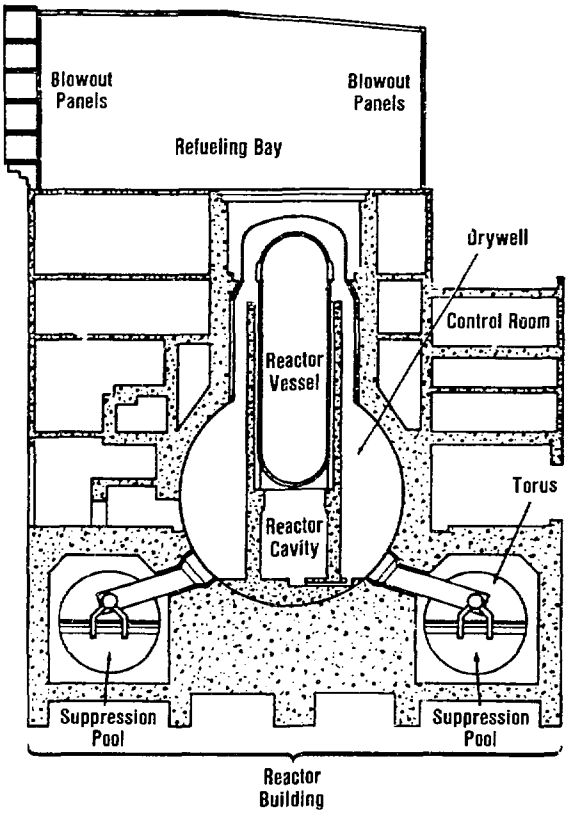

Figure 4. Schematic of the containment design for the Peach Bottom plant.

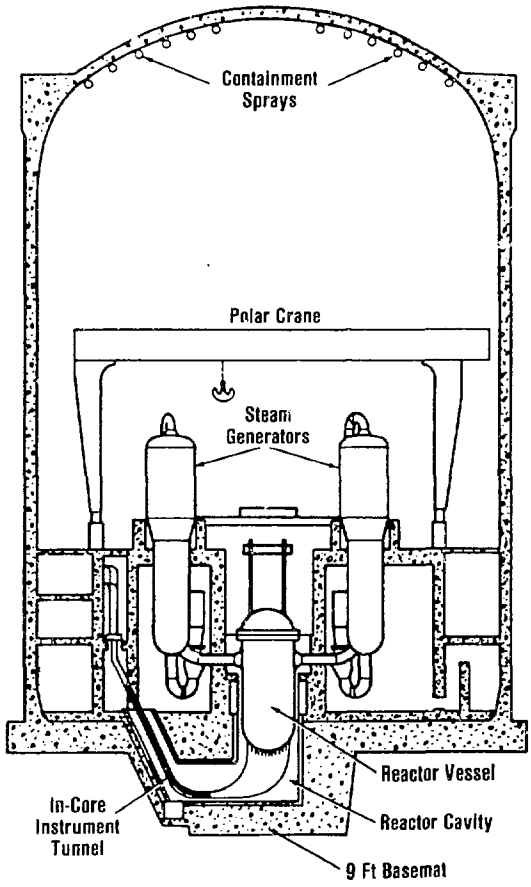

Figure 5. Schematic of the containment design for a typical pressurized water reactor.

levels of radioactivity (above 300 or 400 rem) are lethal and can lead to immediate death. Lower levels can lead to latent cancers and possibly, to genetic damage. But, in going from the first statement of a billion curies to the idea that we want to run a safe plant, in between is the fact that careful engineering principles are used in the design, construction, and operation of the plant to take into account all the possible events that may occur during the lifetime of the plant. These events might occur once a year, or perhaps they are not expected ever to happen. These considerations are built into what are called the engineered safety features of the plant, and these translate into some very expensive equipment representing a large fraction of the total cost of the plant. The cost of nuclear power is basically capital costs, in the construction, design, and fabrication of the plant. The design, operation and construction of the plant set up a series of barriers between the radiation in the reactor core and the public. The fuel matrix itself is a barrier, and around the fuel is cladding, composed typically (in modern reactors) of the zirconium alloy, zircalloy. The reactor vessel is another barrier, and finally, the containment is a barrier. Some people say that emergency planning also forms a barrier to radiological release. Reactors are designed according to very stringent codes unparalleled in any other area of high technology. The NRC regulations for the design, construction and operation of nuclear power plants also are very stringent compared with other areas of technical regulation.

In Figure 6 nuclear risk is compared with some other risks. This shows probability per year of a death for an individual plotted on a logarithmic scale, with 10^{-2} at the very top, going down to 10^{-6}. The probability of an individual dying as the result of living at the site boundary, within a mile or so of the nuclear power plant is 10^{-6}. To put that in perspective, a person who smokes one pack of cigarettes a day has a probability of dying slightly greater than 10^{-3}. There are other examples: the risk of being an engineer is between 10^{-4} and 10^{-5}. The risk from hurricanes ranks just slightly below that from nuclear power. A motorist has a probability of being killed on the road on the order of 10^{-4} The risk of dying as a result of driving a car is arrived at simply by taking our entire population, looking at the number of automobile fatalities per year, and performing a division, which gives us a number on the order of 10^{-4}. For nuclear plants, we have a different situation in that no one has been killed with U.S. commercial nuclear power plants, so we cannot really build up an actuarial database.

How then is this number of 10^{-6} obtained? Since a body count cannot be made the numbers come from analyses of hypothetical accidents. One could postulate certain types of extreme accidents, ways of releasing the large radioactivity from the core into the environment, and to people, and then looking at the subsequent health effects. These have low probabilities; nevertheless, there are some finite probabilities. So we are inferring the risk from nuclear power. Another factor should be pointed out: the risk of living at a site boundary of a nuclear power plant is on the order of 10^{-6}. A few miles out from that point, the risk drops

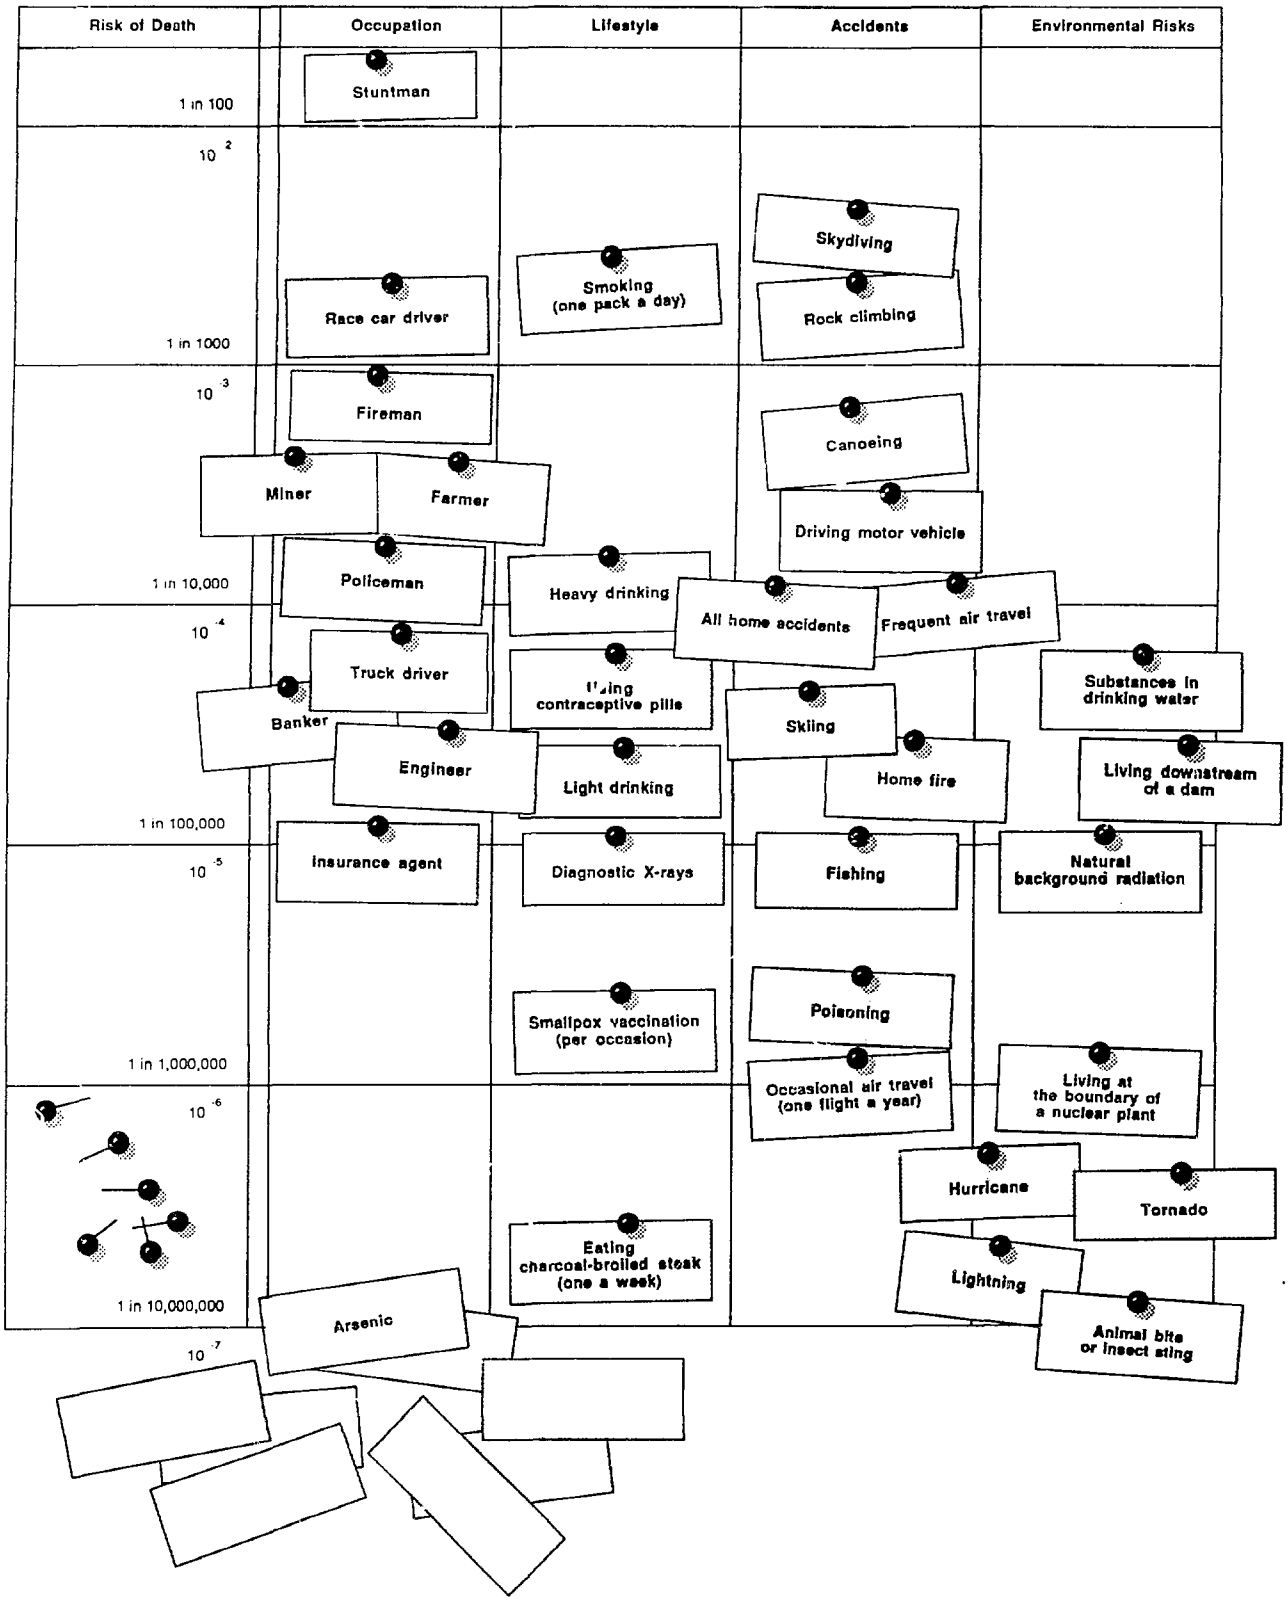

Figure 6. Nuclear risk in perspective.

off very quickly, so that at 7 to 10 miles from the site boundary, the risk very quickly drops by orders of magnitude. How do we calculate these high-consequence low-probability events for nuclear powered plants as opposed to low-consequence high-probability events, which figure in other technologies, such as coal-fired power plants?

In probabilistic risk assessment we are interested in two things: in the probability of an event occurring, and in the consequences. Risk is defined in a symbolic equation, as a probability times consequence (as shown in Figure 7). The equation actually is a complex matrix type of equation, but for our purposes, we are trying to calculate the probabilities of various hypothetical accident sequences and to aggregate

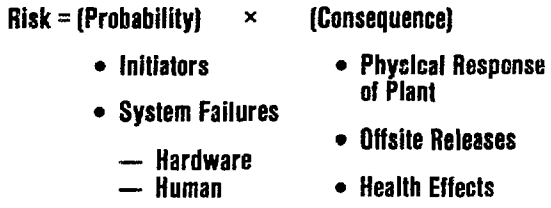

Figure 7. Probabilistic risk assessment. This safety analysis combines probabilistic and deterministic studies to obtain a risk profile.

them in some way to obtain an overall probability. We look for initiators of possible accidents. For example, a break in a pipe in a nuclear power plant, or a large earthquake, or, perhaps, a loss of multiple power systems for the nuclear power plant. These are initiating events. But they are not enough to cause a major accident. The failure of the engineered safety features also has to be considered, because such systems are installed specifically in case of accidents and they mitigate the consequences of these events. The failure of these systems must also become part of the equation. Then, the analysis of these failures ultimately leads to analysis of hardware failures and human failures. On the consequence side of the equation, for each accident sequence postulated or delineated, the physical response of the plant must be examined. This process involves several disciplines including physics, chemistry, and nuclear engineering, to name a few. We must estimate the off-site releases and how they propagate, once the fission products get out of the plant; meteorological considerations must be taken into account. And, finally, we have to look at health effects.

Let me give you a historical perspective for this technology (Figure 8a). It began between 1957 and 1972 and was associated with the early approaches to reactor safety and licensing in the United States. Reactors then were designed according to conservative engineering principles, with a strong containment around the plants, and built-in engineered safety features. We performed deterministic analyses of events, by calculating the temperatures in a reactor or in a containment building, or the forces on pipes as a result of certain events. This approach became the framework for licensing in the United States, and it is basically built into the regulations for nuclear power plant operation. Also in this period, there was informal usage of probabilistic ideas. In fact, in addressing events and how to deal with their consequences, one looked at events that were very likely to occur: events that might occur once a month, once a year, or once during the plant's lifetime. Systems were put into place to accommodate events on such a qualitative scale. Also, probabilistic ideas were introduced in an informal way to look at such questions as the probability of the reactor vessel rupturing. However, the studies were not yet integrated analyses of these probabilistic notions.

On the consequence side, a study was made called WASH 740, which was an analysis of an event that might occur at a nuclear power plant. Basically it was performed to gauge insurance and liability costs for nuclear power plants, and to evaluate how the public is protected by law. The postulate was that a very large amount of radioactivity was released from the building, and the consequences were calculated on that basis. The results showed very large numbers of deaths and health effects from latent cancers. That study lacked a probabilistic perspective. It gave only the consequence side, not the probabilistic side. But by 1972, and probably before that, people were thinking about how to blend the probabilistic notions into a more formal and systematic approach. This culminated in a study called WASH 1400,

sometimes referred to as the Reactor Safety Study (and also the Rasmussen Study).

In this study the Atomic Energy Commission tried to portray the risks from nuclear power plant operation. It was a landmark study in probabilistic risk assessment (Figure 8b). To draw an analogy with the field of physics, this time period for probabilistic risk assessment is roughly analogous to the years 1925 to 1928 for the atomic theory. In a short period of time it laid the foundation for what we call the PRA technology, and the systematic integration of probabilistic and consequential ideas. It also shed light on where the real safety issues were in nuclear power plants. For example, the conventional licensing approach advocated the deterministic analysis of a guillotine pipe rupture in the plant, where one of the large pipes feeding water to the vessel is severed, as if with a hatchet, and then the consequences are calculated in a deterministic way. Using instead probabalistic assessment, WASH 1400 showed that small pipe breaks are the dominant ones in the risk profile.

Transient events, such as loss of power to part of the plant were found to be important. WASH 1400 challenged a notion that one always looks only at single failures in the licensing approach. There is a dictum that one looks at the response of the plant to a transient event in the presence of a single failure. We postulate the additional failure of an active component, and show that the plant can take that event. WASH 1400 showed that is not the limiting case; multiple failures can occur, and, indeed, these have been studied since. The analysis also highlighted the role of the operator of the plant. In looking at failure rates in the plant, it was not the hardware failures for many systems that gave the lead terms in the failure rates, it was such terms as the incorrect performance of a maintenance act, or a failure on the part of an operator to turn a valve to the right position. Another important piece of work in WASH 1400 was the physical analysis of the core-melt sequence, coupling it with an evaluation of the response of the containment. Before that time, people did not seem to recognize the close coupling between how the core melts in the vessel and the responses of the containment. WASH 1400 was the first publication to integrate these factors. Its overall message was that

- **Early approach to reactor safety and licensing**
  - **Conservative engineering principles**
  - **Deterministic analyses**
  - **Current basis for licensing**
- **Informal usage of probabilistic ideas and limited quantification**
- **Wash-740**
  - **Pessimistic consequence analysis**

Figure 8a. A historical perspective on safety analyses, 1957-1972.

- **Wash-1400 (1972 - 1975)**
  - **Accounted for probabilities and consequences**
  - **Established "PRA Technology"**
  - **Adequacy of conventional licensing approach (Small area pipe breaks, transients, single failures)**
  - **Highlighted role of operator**
  - **Core melt/containment analysis**
  - **Risk is low**

Figure 8b. A historical perspective on reactor safety analyses, Wash-1400 to 1975.

- **Transition period (1975 - 1979)**
  - **Lewis report**
  - **Three Mile Island**
- **Post-TMI regulatory requirements and the re-emergence of PRA (1979 - 1981) (Big Rock Point, Limerick, Zion, Indian Point)**

Figure 8c. A historical perspective on reactor safety analysis, Three Mile Island to the present.

the risk was low for nuclear power plant operation (Figure 8b).

WASH 1400 was criticized, extensively analyzed and reviewed. Congress commissioned a report, now called the Lewis Report (after Harold Lewis, University of California, Santa Barbara, who chaired the committee). Their report made certain conclusions about how the Executive Summary of WASH 1400 presented the data, how the uncertainties in the study were not stated correctly or, as they put

it, were "understated." They commented on the scrutability of the report but, overall, they endorsed the methodology used in WASH 1400 and advocated its further usage in the licensing, regulatory and safety areas. Shortly after the Lewis Report, the accident at Three Mile Island occurred (Figure 8c), that led people to ask, Where have we gone wrong? We've had an event which is beyond what we imagined would occur in the light water reactor industry. The regulators reacted to Three Mile Island by imposing many new requirements on nuclear power plants.

A curious thing happened. The people in the probabilistic risk community went back, looked at WASH 1400, and asked, Where did we go wrong in our analysis? We did a study of risk, and we seem to have missed Three Mile Island in it somewhere. But, upon closer inspection, it was, in principle, in there. The initial study done for WASH 1400 was based on two specific plants. One was a Westinghouse plant, the Surry Plant, and the other was at Peach Bottom, a General Electric Plant. It was thought at the time that two power plants, one a BWR and one a PWR, were fairly representative of the hundred or more nuclear plants that would be in place by the last part of the 20th century. We were wrong; each plant is unique. We have assessed risk, only if we make a risk assessment for each individual power plant, delineating the dominant contributors to risk at each plant. If we go back and look at the Three Mile Island plant, which is a Babcock-Wilcox plant with certain design and operational procedures, we find that the particular data and quantification in WASH 1400 is not representative of that plant. Qualitatively however, the event, with many other sequences, was delineated in WASH 1400 and by using failure data and plant modeling appropriate to TMI, the event could be explained within the framework of WASH 1400. This provided an impetus for a harder look into the area of probabilistic risk assessment by the regulators and also by the nuclear industry.

Industry itself took the big initiative in probabilistic risk assessment. They performed full plant-specific, probabilistic risk assessments. One of the first studies was made for the Big Rock Point Plant, located out in the Midwest. This small plant is rather old, producing about 300 MW of electricity. A lot of requirements were put on it since the Three Mile Island incident, and, looked at from the risk perspective, it did not make sense to do the types of things that the regulators were promoting on a deterministic basis. Following that, full-scale studies were done for three other plants, which had the special consideration that they were in areas of high population density: Indian Point, for example, is about 36 miles from Times Square, New York, Zion is near Chicago, and Limerick is close to Philadelphia. The basic conclusions, based on plant-specific features, were that the risks were low. Some features of each plant were identified as the risk outliers, and they were correctable. For example, Unit Two, at Indian Point, was found to be vulnerable to a seismic event, because two buildings would bounce against each other. This was very simple to fix: a bumper was placed between the two buildings. The control room ceiling of Unit Three at Indian Point was vulnerable to collapse, and it was reinforced appropriately. This brings us to the present. About 20 additional probabilistic risk assessments have been done. We at Brookhaven have participated in this activity, partly by performing these assessments, and partly by providing peer reviews for the Nuclear Regulatory Commission.

Peer review, in this case, is not a simple reading of a report for two or three days, doing some "back-of-the-envelope" calculations, and then giving a referee's assessment. It is more of a six to ten man-year effort over a period of one to two years, that includes detailed exchange of information with utilities and their representatives and participation in public hearings. Very large computer programs are involved, as well as replication of very detailed calculations, promulgation of new calculations, finding new events that might occur, and, ultimately, providing a reassessment of the risk.

The PRA methodology has several aspects (Figure 9). One is the logic model, which identifies events that are the so-called "bad actors" in the risk arena. In the way the logic models go, we start with two types of logic trees, with their foundation in Boolean algebra. One is the event tree approach, which is inductive. It moves forward in time to delineate events through a two-level logic type trees; yes/no, fail/success. The other is a fault-tree approach, where one starts with a top event, which is the undesired event, and goes backward in time to

- **Logic models**
  - **Event trees (inductive)**
  - **Fault trees (deductive)**
- **Data (including judgment)**
- **Physical models**
- **Health and economic loss models**

Figure 9. Methodology of probabilistic risk assessment.

find out what has led to this event. Current PRA combines both approaches, to go simultaneously backward and forward in time. The principal advantage is that it reduces the combination of events that would be present if one or the other type approach were used alone.

Another very important feature of probabilistic risk assessment concerns the source of the data used to quantify risk assessments. The data is obtained in part from the nuclear power plant experience itself, and partly from other industries. For example, we look at valve failures in fossil fuel plants and ask rather sophisticated questions of how the data base from another part of the sample space applies to the events that we want to quantify. Another part of the data base is judgment. Sometimes events are of very low probability or of very high uncertainty, so that judgment has to supply the data there. There is nothing wrong with this; in fact, the Bayesian point of view in probabilistic theory states that judgment is our sole source of data or, at least, a starting point. The physical models describe how the core melts down, how the containment behaves, and how off-site consequences progress. Then we assess the final damage, the health effects, and the economic losses.

Let me give a clearer picture of the logic models using a highly simplified event tree. In Figure 10a, we start with an initiating event. Suppose the hypothetical machine we are considering is made up of two systems, A and B. Given that Event I occurs, we ask then whether System A is found in a success state or in a failure state. We follow the appropriate branch, and ultimately put probabilities on those branches. The same process is followed for System B. Is it successful or not? If there are N systems, there will be 2^N possible states at the end. Then we form a Boolean expression with I and A and B. Looking at those N states we can ask whether there will be trouble or not. For example, given an undesired initiating event in a reactor, and if both safety systems succeed, we would expect not to have trouble, i.e., the accident is mitigated.

When only System A or System B is sufficient to perform the overall safety function for Event I, then an event in which A is successful, but B fails will be a safe event. Similarly, the complementary situation where A fails and B succeeds will not be hazardous. The bar in Figure 10a denotes the event does not occur. The figure shows three events that present no hazard. But if the failure of both systems occurs, then there is core melt. The fault trees enter into the picture when we look at System A and find that it is made up of sub-systems. We go backward in time through the fault tree again using "AND" gates and "OR" gates in the basic logic models.

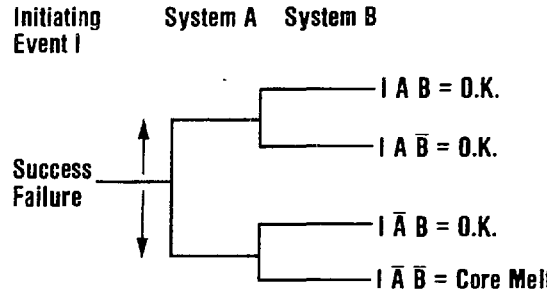

Figure 10a. A highly simplified event tree. Note that the bar denotes that the event does not occur.

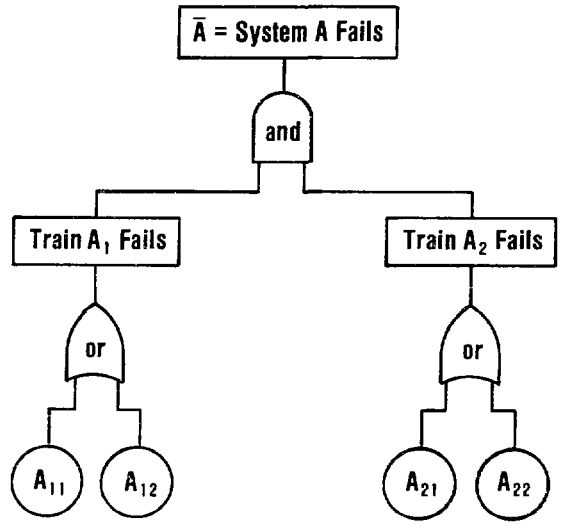

Figure 10b. A simplified fault tree.

For those familiar with Boolean algebra, it follows that if A_1 fails and A_2 fails, then the overall System A fails. We assume here that System A is now made up of two trains, both of which have to fail in order to bring about the failure of System A (Figure 10b). The real reactor situation is more complex than this, with very large full event trees; these are put together, and finally quantified via probabilistic expressions. This analysis requires people who are familiar with the procedure of fault tree manipulation and quantification. However, an analyst or a probabilistic theorist could not alone do a good job of developing a fault. A Plant System Analyst who knows the plant is needed; preferably the original designer of the plant should participate very actively in this process. Also, an engineer familiar with the hardware should participate, and people familiar with human factors should look at A_{12}, if here the operator has failed to open the valve in time. Constructing fault trees is a complicated process that requires many people, who ultimately have to know the plant well.

The event tree is an inductive process, and in constructing one we are only as good as our imagination. Ultimately we face the difficult problem of completeness — of whether we found all the significant events. What has happened so far? In terms of the prediction of a core meltdown, or the frequency of a core meltdown, many initiators have been included, for example system failures, earthquakes, in-plant fires, floods (Figure 11). Studies have identified a very large class of initiators for the event trees (i.e., Event I on the event tree). The combined event tree/fault tree approach is used to define the endpoint, the core melt accident sequences, and then, when we do the quantification, we try to account for uncertainties in data, modeling the success criteria that

- **Many initiators included (system failures, earthquakes, etc.)**
- **Combined fault tree/event tree approach**
- **Quantification accounts for uncertainties**
- **Current results for ~20 plants yield $P_{CM} \sim 10^{-3} - 10^{-5}$/year**

Figure 11. Prediction of core meltdown frequency.

we might put into a system. Probablistic risk assessments have been made for about 20 plants and they predict a mean value for core-melt frequency per year of between 10^{-3} and 10^{-5}.

Turning next to the physical analysis we will follow the steps leading to a hypothetical core meltdown event (Figure 12). For reference, the containment building is 200 feet high. The progression starts with some initiating events in the reactor vessel perhaps a pipe break. This leads to a loss of water in the vessel, ultimately uncovering the core. Once the core is uncovered, it heats up. It loses its immediate heat sink, and if the heat sink is not restored, the core will melt.

The core is shown schematically in Figure 12 by the dotted line. As the core melts, the debris drips down toward the lower regions of the vessel. Failure of the vessel occurs through the lower head, and then ultimately there is interaction between the core and the concrete. Note that the bottom of the containment is made up of about nine feet of concrete, and the core debris is very hot at this point so that it will start eroding the concrete; in some scenarios it may be 2000°F, or even as high as 4000°F. What does this do to the containment? The thermal hydraulic loading of the containment is shown on the left of Figure 12. This will be in the form of steam, that comes at the initiation of the event. Since the reactor water is at very high pressure in the vessel, it comes out in flashes of steam, which pose a thermal-hydraulic load on the containment. Non-condensibles also may come out. As a result of the uncovering of the core, the zirconium cladding reacts with steam producing an exothermic reaction. Hydrogen is also produced, which is a non-condensible, but even worse, it is a combustible. As a result of this process there may be combustion loading on the containment, also there may be direct loading by the aerosols that are emitted. This could lead to a loading in the containment, due to thermal and chemical effects. One is ultimately interested in the behavior of the fission products. That is where the risk comes from, as shown on the right-hand side of Figure 12. Aerosols are produced, initially in the vessel as a result of the core degradation and melting.

If the fission products stay within the vessel long enough, possibly they could settle out and

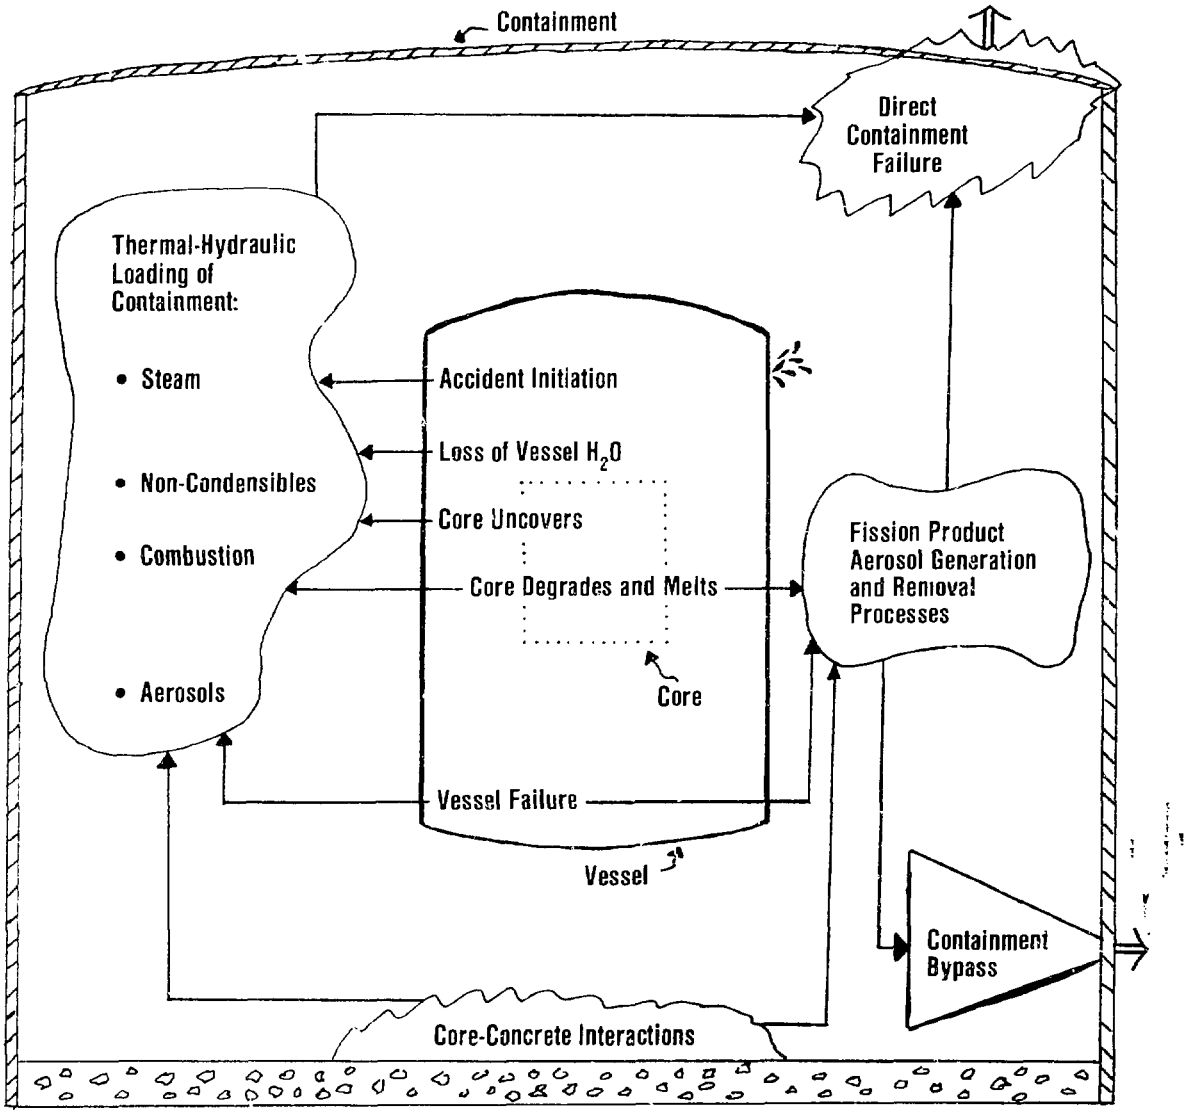

Figure 12. A hypothetical core meltdown event.
The core is shown by the dotted line.

immediately plate out onto relatively colder regions of the vessel and not go any further. If they escape from the vessel, then one asks about the processes of removal of fission products in the containment building itself. If there is agglomeration of aerosols in the building this could lead to plating out, again on to structures. The containment building is complicated and has many structures that provide plating-out areas. Then what happens? As shown in Figure 12 the fission products may still be available to go off site, when thermal hydraulic loading leads to a direct containment failure because the pressures were too high. Either an overall static pressure due to a gradual steam loading, or an abrupt dynamic pressure due to hydrogen deflagration, or even a detonation, could rip open the containment building or cause a slow leak in the containment building. Whatever surviving fission products were in the containment atmosphere could then be released. This is called the source term, that is, the radiation that comes out of the containment building into the atmosphere. We specify the type of radioactivity, the energy with which it comes out, and the location at which it comes out. This is not the whole story. There is another possibility, in

which we actually bypass the containment, either as a result of an undetected opening anywhere in the containment or by a combination of system failures. The containment is designed to be essentially airtight; the design basis leak rate is a half percent per volume of exchange per day.

If the containment is violated, initially, or, as a result of the accident scenario, then there would be a direct route to the exterior without making an assault on the overall wall structure of the containment building. That is another way of getting the source term out of the building.

We are much concerned with source term analysis at Brookhaven. Part of our interest is connected with the fact that the quantity of fission products that were emitted during the Three Mile Island accident were much less than expected, particularly for the iodine species. Only a very small amount was released, and basically this was due to the fact that iodine did not come out in the elemental phase, but formed compounds with cesium; cesium hydroxide also was formed in the building, and these compounds behaved as aerosols and tended to plate out, which led to a much reduced source term coming from these accidents. There are many imponderables associated with accidents at nuclear plants. We are entering into a realm that has not been subjected to large-scale full experiments: we only have limited experiments with simulants or large-scale prototypes, so we have to rely on our analytical tools. In doing this, we are faced with large uncertainties. Part of what we are doing is trying to get a better handle on these uncertainties to quantify them in a better way, and trying to understand which accidents are the controlling ones as far as risk goes. The basic perspective that has come out of our work and from contemporary PRAs is that most core meltdown accidents do not destroy the containment.

These findings are contrary to what WASH 1400 assumed. WASH 1400 predicted a rather high probability of containment failure with a core meltdown. Subsequent analyses in the period 1980-1985 show that probably 99% of core melt events do not lead to containment failure. Furthermore, a core meltdown event typically takes many hours to develop. The event shown schematically in Figure 12 might take 10 to 20 hours from the time of its initiation to the time of failure of the containment. The failure of the vessel itself could take a quarter of a day or so, depending on the scenario. Furthermore, the containment holds in fission products well, even in those cases where the engineered safety features that are put into the containment happen to fail. There are cooling systems for reducing the temperature of the containment, and also fission product removal systems in the atmosphere; both are engineered features put in by design. Even if these fail, the containment, in a passive way, allows the fission products to settle particularly during events where containment does not fail very quickly.

At present we are looking at severe fast-moving accidents that can occur in two or three hours and which might bypass the containment. Events that challenge the containment by slow pressurization give a very low risk. The event that somehow finds its way out of the containment through some bypass routes, or which causes the containment to fail very quickly seems to be the dominant term in the risk. As we go further into this analysis, we are narrowing the arena of high-risk events, focussing on the third item bullet in Figure 13.

Let me conclude with some further directions and applications. Probabilistic risk assessment has given us improved descriptions of accident sequences and consequences that

- **Most core meltdown events do not fail containment (and take several hours to develop)**
- **Containment allows for passive removal (settling) of fission product aerosols**
- **Current risk focus is on low probability fast moving accidents which bypass containment**

Figure 13. Perspective in the analysis of containment.

are of interest to those who regulate nuclear power, and to plant operators who have found that such studies gives them an integrated sense of how a plant may operate during an abnormal phase. The studies give a sense of why the engineered safety features are there, how they are coupled together, and where dependencies between systems could lead to

multiple failures — or what are called common mode-type failures. The work has focussed on the need for maintenance of reactors, highlighting the areas of the reactor that need particular attention to maintain safety. How do we use risk assessment to assure operational safety or how do we make sure, once we do a probabilistic risk assessment and find that the risk is low, that this is true for the 30-year lifetime of the plant? What should we do on a day-to-day basis to make sure that the prediction comes true?

We have a large program at Brookhaven that addresses these questions. In doing the analyses, we identified vulnerabilities in the plant. For example, about a year ago we made a study called the Systems Interaction Study, that combined the fault-tree and event-tree approach to find a systems interaction in the electrical system at the Indian Point Plant. Our results led the operator of the plant to alter the design before being told to do so by the Nuclear Regulatory Commission. The conventional design basis review process used by the Nuclear Regulatory Commission failed to reveal this event; however, the study at Brookhaven uncovered it and led to its correction and a recognition of its importance to the safe operation of the plant.

An overall perspective of risk is very helpful in making decisions connected with regulations and licensing. It gives the regulator an extra view of the safety of the plant, beyond the conventional approach, and so it is coming more and more to the front in the licensing arena. The big drawback is that the approach is rather complicated, and the results often have to be represented in a complex fashion. The results must be conveyed to the regulators in a clear unambiguous way. Furthermore, the regulatory process in the United States involves not only the Nuclear Regulatory Commission interacting with potential licensees, but it is really an adjudicatory process carried out in the courts of law. For us, it is a real challenge to attend the hearings and, under expert testimony, to convey the thought processes underlying some of these analyses. Risk assessment has been helpful in setting priorities for safety research.

Lastly, a PRA is a communication tool. By presenting bottom-line numbers to the public and their elected representatives, we can give a perspective on where we think the risk is. How well we succeed is another matter, but this is a way of quantifying, in a rational manner in an irrational world, where we stand.

U.S.Government Printing Office: 1987-711-035/40012

ABOUT THE AUTHOR

Robert Bari is a senior physicist and an associate departmental chairman in the Department of Nuclear Energy. He is responsible for the engineering and risk assessment programs that are sponsored by the U.S. Nuclear Regulatory Commission. His research interests include reliability analysis and core meltdown analysis. Dr. Bari received his bachelor's degree from Rutger's University and his Ph.D in physics from Brandeis University. Before joining the Department of Nuclear Energy in 1974, he worked in solid state physics at Brookhaven, at SUNY at Stony Brook, and at the MIT Lincoln Laboratory.